Neural Multi-Step Reasoning for Question Answering on Semi-Structured Tables

Till Haug
Veezoo AG
Zurich, Switzerland
till@veezoo.com

Octavian-Eugen Ganea
ETH Zurich
Zurich, Switzerland
octavian.ganea@inf.ethz.ch

Paulina Grnarova
ETH Zurich
Zurich, Switzerland
paulina.grnarova@inf.ethz.ch

arXiv:1702.06589v1 [cs.CL] 21 Feb 2017

Abstract

Advances in natural language processing tasks have gained momentum in recent years due to the increasingly popular neural network methods. In this paper, we explore deep learning techniques for answering multi-step reasoning questions that operate on semi-structured tables. Challenges here arise from the level of logical compositionality expressed by questions, as well as the domain openness. Our approach is weakly supervised, trained on question-answer-table triples without requiring intermediate strong supervision. It performs two phases: first, machine understandable *logical forms* (programs) are generated from natural language questions following the work of [Pasupat and Liang, 2015]. Second, paraphrases of logical forms and questions are embedded in a jointly learned vector space using word and character convolutional neural networks. A neural scoring function is further used to rank and retrieve the most probable logical form (interpretation) of a question. Our best single model achieves 34.8% accuracy on the WikiTableQuestions dataset, while the best ensemble of our models pushes the state-of-the-art score on this task to 38.7%, thus slightly surpassing both the engineered feature scoring baseline, as well as the Neural Programmer model of [Neelakantan *et al.*, 2016].

1 Introduction

Teaching computers to answer complex questions expressed in natural language is an old longing of artificial intelligence that requires sophisticated reasoning and human language understanding. In this work, we investigate generic natural language interfaces for semi-structured databases. Typical questions for this task are topic independent and require multi-step reasoning involving discrete operations such as aggregation, comparison, superlatives or arithmetic. One example is the question `What was the difference in average attendance between 2010 and 2001?` which has to be answered based on a table containing information from soccer games[1]. This line of research is practically relevant for automated systems that support interactions between non-expert users and databases without requiring specific programming knowledge.

Question-Answering (QA) systems are often faced with a trade-off between the openness of the domain and the depth of logical compositionality hidden in questions. One example are systems able to answer complex questions about a specific topic (e.g. [Wang *et al.*, 2015]). Unsurprisingly, these systems often struggle to generalize to other, more open domains. On the other side, topic-independent QA systems that can potentially interrogate large databases are usually limited to simple look-up operations (e.g. [Bordes *et al.*, 2014a]).

Here, we propose a novel weakly supervised model for natural language interfaces operating on semi-structured tables. Our deep learning approach eliminates the need for expensive feature engineering in the candidate scoring phase, while being able to generalize well to never-seen before data. Each natural language question is translated into a set of computer understandable candidate representations, called logical forms, based on the work of [Pasupat and Liang, 2015]. Further, the most likely such program is selected in two steps: i) using a simple algorithm, logical forms are transformed back into *paraphrases* (textual representations) understandable by non-expert users, ii) next, these raw strings are further embedded together with the respective questions in a jointly learned vector space using convolutional neural networks over character and word embeddings. Multi-layer neural networks and bilinear mappings are employed as effective similarity measures and combined to score the candidate interpretations. Finally, the highest ranked logical form is executed against the input data to retrieve the answer that will be shown to the user. Our method uses only weak-supervision from question-answer-table input triples, without requiring expensive annotations of logical forms or latent operation sequences.

We empirically confirm our approach on a series of experiments on WikiTableQuestions [Pasupat and Liang, 2015], a real-world dataset containing 22,033 pairs of questions and their corresponding manually retrieved answers with about 2,108 randomly selected Wikipedia tables. The inherent chal-

[1]Taken from the WikiTableQuestions dataset: `http://nlp.stanford.edu/software/sempre/wikitable/viewer/#204-590`

lenges of this dataset include i) the small number of training examples, ii) the complexity of questions that generally require compositionality over multiple simpler operations, iii) generalization to completely unseen tables and domains at test time, and iv) lack of strong supervision. An ensemble of our best models is able to reach state-of-the-art accuracy of 38.7% on this task, showing that neural networks are promising at outperforming engineered systems when dealing with complex questions. We provide further insight into our method by comparing against a set of baselines (including engineered feature systems) and model variations, as well as showing ablation and error analysis studies.

2 Related Work

We now briefly highlight prior research relevant to our work. There are two main types of QA systems: semantic parsing-based and embedding-based.

Semantic parsing-based methods perform a functional parse of the question that is further converted to a machine understandable program and executed on a knowledgebase or database. A big obstacle for semantic parsers is the need for annotated logical forms when dealing with new domains. To tackle this problem, our method follows recent work of [Reddy *et al.*, 2014; Kwiatkowski *et al.*, 2013] that relies solely on weak-supervision through question-answer-table triples. In the context of QA for semi-structured tables and dealing with multi-compositional queries, [Pasupat and Liang, 2015] generate and rank candidate logical forms with a log-linear model trained on question-answer pairs. The logical form with the highest model probability is then considered as the correct interpretation and is executed. In this work, we generate logical form candidates in the same way as [Pasupat and Liang, 2015]. While they resort to hand-crafted features to determine the relevance of the candidates for a question, we use automatic learning of representations, thus benefiting from generalization and flexibility. To do so, we embed each question and the paraphrases of the respective candidate logical forms into the same vector space, making use of similarity metrics for scoring. Paraphrases have been successfully used to facilitate semantic parsers [Wang *et al.*, 2015; Berant and Liang, 2014]. While [Berant and Liang, 2014] is suited for factoid questions with a modest amount of compositionality, [Wang *et al.*, 2015] targets more complicated questions. Both of these paraphrase-driven QA systems differ from our work as their scoring relies on hand-crafted features.

[Neelakantan *et al.*, 2016] also focus on compositional questions, but instead of generating and ranking multiple logical forms, they propose a model that directly constructs a logical form from an embedding of the question. A list of discrete operations are manually defined and each operation is parametrized by a real-valued vector that is learned during training. A separate recurrent neural network (RNN) is used for modeling the history of selected operations and the question representation. The probability distributions over operations and columns are induced using the question embedding, the history and an attention vector. While their model performs well on semi-structured tables, understanding how it generated the logical form is not trivial. Recently, [Yin *et al.*, 2015] propose Neural Enquirer, a fully neural, end-to-end differentiable network that executes queries across multiple tables. They use a synthetic dataset to demonstrate the abilities of the model to deal with compositionality in the questions.

Embedding-based methods represent the question and the answer as semantic vectors. Compatibility between a question-answer pair is then determined by their similarity in the shared vector space. Existing approaches often represent questions and knowledgebase constitutes in a single vector using simple bag-of-words (BOW) models [Bordes *et al.*, 2014a; Bordes *et al.*, 2014b] under the framework of memory networks. [Dong *et al.*, 2015] propose a multi-column convolutional neural network to account for the word order and higher order n-grams. Even though we embed questions along paraphrases instead of answers, our method relates to embedding-based models since the key challenge is learning representations. However, whereas these models are applied on datasets that require little compositional reasoning, our work targets questions whose answers ask for multi-step complex deductions and operate on semi-structured tables instead of structured knowledgebases. Representation learning using deep learning architectures has been widely explored in other domains, e.g. in the context of sentiment classification, [Kim, 2014; Socher *et al.*, 2013], or for image-hashtag prediction [Denton *et al.*, 2015].

QA systems also differ in the knowledge structure they reason on, which can impose additional challenges. Systems vary from operating on structured knowledge bases [Bordes *et al.*, 2014b]; [Bordes *et al.*, 2014a] to semi-structured tables [Pasupat and Liang, 2015], [Neelakantan *et al.*, 2016], [Jauhar *et al.*, 2016] and completely unstructured text, which is related to information extraction [Clark *et al.*, 2016]. We focus on semi-structured tables that face the trade-off between degree of structure and ubiquity.

3 Model

We now proceed with the detailed description of our QA system. In a nutshell, our model runs through the following stages. For every question q : i) a set of candidate logical forms $\{z_i\}_{i \in I_q}$ is generated using the method of [Pasupat and Liang, 2015]; ii) each such candidate program z_i is paraphrased in a textual representation t_i that offers accuracy gain, interpretability and comprehensibility ; iii) all textual forms t_i are scored against the input question q using a neural network model; iv) the logical form z_i^* corresponding to the highest ranked t_i^* is selected as the machine-understandable translation of question q; v) z_i^* is executed on the input table and its answer is returned to the user. Our contributions are the novel models that perform the steps ii) and iii), while for i), iv) and v) we rely on the work of [Pasupat and Liang, 2015] (henceforth: PL2015).

We detail each of these steps in the subsequent sections.

3.1 Candidate Logical Form Generation

We use the method of PL2015 to generate a set of candidate logical forms with respect to a question. Specifically, given a pair of input table and question to be answered, the table is first transformed into a knowledge graph (KG). Next,

Algorithm 1 Recursive paraphrasing of a Lambda DCS logical form. The + operation is defined as string concatenation with spaces. Details about Lambda DCS language can be found in [Liang, 2013].

procedure PARAPHRASE(z) ▷ z is the root of a Lambda DCS logical form
  switch z **do**
    case Aggregation ▷ e.g. count, max, min...
      $t \leftarrow$ AGGREGATION(z) + PARAPHRASE($z.child$)
    case Join ▷ join on relations, e.g. λx.Country(x, Australia)
      $t \leftarrow$ PARAPHRASE($z.relation$) + PARAPHRASE($z.child$)
    case Reverse ▷ reverses a binary relation
      $t \leftarrow$ PARAPHRASE($z.child$)
    case LambdaFormula ▷ lambda expression $\lambda x.[...]$
      $t \leftarrow$ PARAPHRASE($z.body$)
    case Arithmetic **or** Merge ▷ e.g. plus, minus, union...
      $t \leftarrow$ PARAPHRASE($z.left$) + OPERATION(z) + PARAPHRASE($z.right$)
    case Superlative ▷ e.g. argmax(x, value)
      $t \leftarrow$ OPERATION(z) + PARAPHRASE($z.value$) + PARAPHRASE($z.relation$)
    case Value ▷ i.e. constants
      $t \leftarrow z.value$
  return t ▷ t is the textual paraphrase of the Lambda DCS logical form
end procedure

information from the KG facilitates the process of parsing a question into a set of candidate logical forms. This is done using a semantic parser that recursively builds up logical forms by repeatedly applying deduction rules. Each candidate logical form is represented in Lambda DCS form [Liang, 2013] and can be transformed into a SPARQL query, whose execution against the KG yields an answer.

For instance, the question `How many people attended the last Rolling Stones concert?` will be translated into a set of candidate logical forms, among which the correct one is: **R**[λx[`Attendance.Number.`x]] `.argmax(Act.RollingStones,Index).`

3.2 Converting Logical Forms to Text

We describe our proposed paraphrasing algorithm to transform logical forms into textual representations. These, and not the original logical forms, are further scored against the input question. Besides the advantages of interpretability and comprehensibility, we also observe a quality gain when using paraphrases in comparison with ranking directly based on the string representations of the original Lambda DCS expressions (details in section 4.4).

Usability and extensibility of QA systems may benefit from revealing the translation of the human-language question in machine-language (e.g. Lambda DCS in our case). This is called *transparency*, i.e. the property of revealing the generated program from an input question in its raw format. However, this might not achieve *comprehensibility*, the characteristic of a system to be understandable by non-technical users. Paraphrasing the logical form z into a textual representation t satisfies both of these properties as it yields a non-expert understandable description of the executed program based on an input question.

Given a logical form z in Lambda DCS, we paraphrase it using a simple algorithm that recursively traverses the tree representation of z starting at the root. The translation operations associated with each node type can be seen in Algorithm 1. As an example of how this algorithm works, the correct candidate logical form for the question mentioned in section 3.1, namely `How many people attended the last Rolling Stones concert?` will be mapped to the paraphrase `Attendance as number of last row where act is Rolling Stones.`

3.3 Joint Embedding Model

For each question q we generate a set of logical forms z_i and apply Algorithm 1 to retrieve their corresponding paraphrases t_i. Subsequently, questions and paraphrases are embedded in a jointly learned vector space. Each t_i will be scored based on the similarity with question q defined with a neural network acting on top of their corresponding embeddings. Features used by our scoring system are learned automatically without the need for hand-engineering. We use two separate convolutional neural networks (CNNs) for question and paraphrase embeddings, on top of which a max-pooling operation is applied. The CNNs receive as input token embeddings consisting of concatenation of word and character vectors. The details of our model are outlined in the following sections. For readability, some hyper-parameter values are shown in section 4.2.

Token Embedding

We now detail the neural architecture used to embed tokens of an input piece of text (e.g. question, paraphrase). This model is depicted in Figure 1. Every token in our vocabulary is parametrized by a word and a character embedding, both learned during training.

We find that word vector initialization affects both the convergence speed and quality of our method. A typical choice for this initialization is to use pre-trained vectors learned from

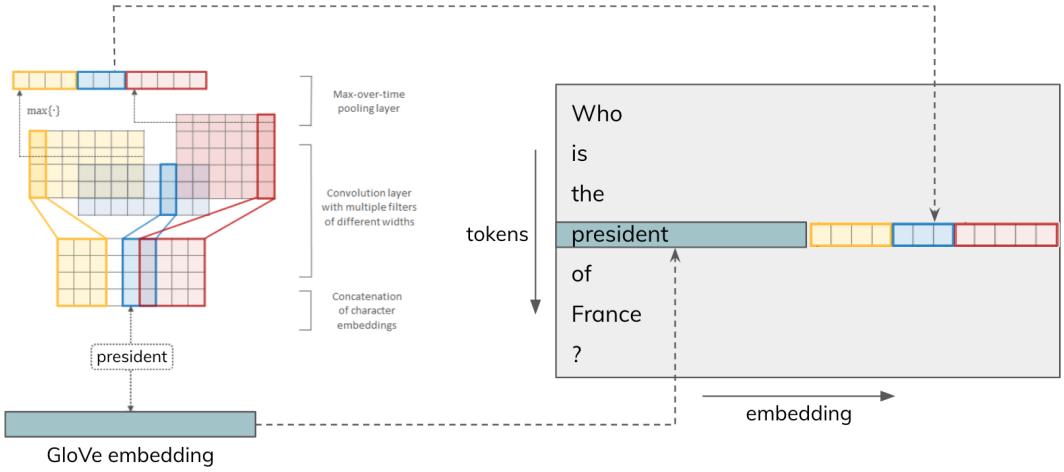

Figure 1: Proposed architecture to convert a sentence into a token embedding matrix. Each token vector is a concatenation of the corresponding word and character embeddings. The latter is obtained using a character CNN consisting of filters of different character widths (depicted of size 2, 3 and 4). Parts of figure taken from [Kim *et al.*, 2015].

unsupervised textual data. These models are known to encode both syntactic and semantic information, e.g. similar or related words like synonyms are mapped to close points in the vector space. We experiment with two different popular methods, namely GloVe [Pennington *et al.*, 2014] and Word2vec [Mikolov *et al.*, 2013], comparing them also with random initializations. Embeddings for tokens not in the vocabulary are randomly initialized by sampling each component from a uniform distribution $U[-0.25, 0.25]$.

When comparing the textual representations of the logical forms to the input question, we notice that in many cases similar numbers or dates are written in different ways, e.g. `1,000,000` vs `1000000`. Other common sources of noise are rare words and word misspellings. These are different or unknown tokens in the vocabulary for which word embeddings alone would perform poorly. One technique to mitigate these issues inspired from [Kim *et al.*, 2015; Zhang *et al.*, 2015] is to use character embeddings in addition to word vectors. Token vectors are then obtained using a CNN over the constituent characters. Our CNN model uses multiple filter widths, followed by a max-over-time pooling layer and concatenation with the respective word vector.

Sentence Embedding

We map both the question q and the paraphrase t into a joint vector space using sentence embeddings obtained from two jointly trained CNNs. The input to the CNNs are tensors of shape $\mathbb{R}^{s \times d}$, containing an embedding for each token, where s is the maximum input length and d is the dimension of the token embedding. We use filters spanning a varying amount of tokens with the widths from a set L. For each filter width $l \in L$, we learn n different filters, each of dimension $\mathbb{R}^{l \times d}$. After the convolution layer, we apply a max-over-time pooling on the resulting feature matrices which yields, per filter-width, a vector of dimension n. Next, we concatenate the resulting max-over-time pooling vectors of the different filter-widths in L to form our sentence embedding. The final embedding size produced by this architecture is thus $n|L|$. As non-linearity we use Exponential Linear Units (ELUs) [Clevert *et al.*, 2015].

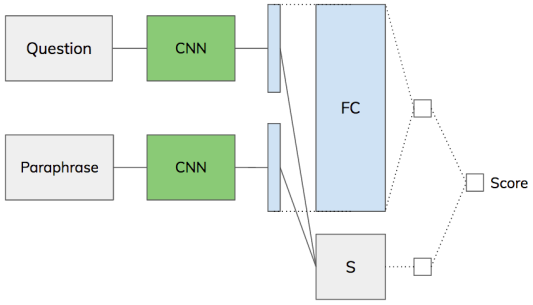

Figure 2: Design of our best single model, CNN-FC-BILINEAR-SEP. Embedding the question and paraphrase using two CNNs yields two vectors. The final score for the paraphrase is a linear combination between the output of two networks: a bilinear mapping and a fully connected network.

3.4 Neural Similarity Measures

We denote the sentence embedding of the question q by $u \in \mathbb{R}^k$ and of the paraphrase t by $v \in \mathbb{R}^k$, respectively. We experiment with various neural-based similarity scores between u and v denoted as follows:

1. $u^T v$ (DOTPRODUCT)
2. $u^T S v$, with $S \in \mathbb{R}^{k \times k}$ a parameter matrix learned during training. (BILINEAR)
3. (u, v) concatenated followed by two sequential fully connected layers. (FC)
4. BILINEAR concatenated with u and v and followed by fully connected layers. (FC-BILINEAR)

5. BILINEAR and FC linearly combined with learned weights (Figure 2). (FC-BILINEAR-SEP)

The best performing model is FC-BILINEAR-SEP as shown in Table 1. Intuitively, BILINEAR and FC are able to extract different interaction features between the two input vectors, while their linear combination retains the best of both models.

Training Algorithm

For training, we define two sets $\mathcal{P}$ (positive) and $\mathcal{N}$ (negative) that contain pairs (q, t) of questions and paraphrases of logical forms that, when executed on their respective tables, give the correct or incorrect answer[2], respectively. Our models presented in Section 3.4 map such a pair to a real valued score representing question-logical form similarity, denoted here by the function $\Phi : Q \times T \rightarrow \mathbb{R}$. We use a max-margin loss function (with margin θ) aiming to rank pairs in $\mathcal{P}$ above pairs in $\mathcal{N}$:

$$L(\mathcal{P}, \mathcal{N}) = \frac{1}{|\mathcal{P}||\mathcal{N}|} \sum_{p \in \mathcal{P}} \sum_{n \in \mathcal{N}} \max(0, \theta - \Phi(p) + \Phi(n))$$

We also experiment with cross entropy loss function, but achieve significantly worse results.

4 Experiments

4.1 Dataset

We use the train-validation-test split of WikiTableQuestions dataset containing 9,659, 1,200 and 4,344 questions, respectively. We obtain about 3.8 million training (q, t, l) triples from PL2015, where l is a binary indicator of correctness (whether the logical form gives the correct answer when executed). During training we ignore questions for which a single matching pair (q, t) is not present. The percentage of questions for which a candidate logical form exists that evaluates to the correct answer is called *oracle* score. PL2015 report an oracle score of 76.7%, but a manual annotation by [Pasupat and Liang, 2015] reveals that PL2015 can answer only 53.5% of the questions correctly. The difference can be explained by incorrect logical forms that give the correct answer by chance.

4.2 Training Details

The neural network models are implemented using TensorFlow [Abadi *et al.*, 2016] and trained on a single Tesla P100 GPU. Training takes approximately 6 hours, while 50'000 mini-batches are processed. Generating the textual representations from logical forms for all the questions with PL2015 takes about 14 hours on a 2016 Macbook Pro computer.

The vocabulary contains 14,151 tokens. We obtain the best results when initialising the word embeddings with the 200 dimensional GloVe vectors. Using higher dimensional vectors does not result in significant gains in accuracy. The size of the word embeddings $x_{emb} = (x_{glove}, x_{char})$ is set to $d = d_{glove} + d_{char} = 200 + 192 = 392$. The sentence embedding CNNs span multiple tokens with widths of $L = \{2, 4, 6, 8\}$, while for the character CNN we use widths spanning 1, 2 and 3 characters. The two fully connected layers in the FC models have 500 hidden neurons, which we regularize using dropout [Srivastava *et al.*, 2014] with a keep probability of 0.8. We use a mini-batch size of 100, each batch containing 50 different questions q with one positive t_1 and one negative t_2 paraphrase. We set the margin θ of the loss function to 0.2. Loss minimization is done using the Adam optimizer [Kingma and Ba, 2014] with a learning rate of $7 * 10^{-4}$. All hyperparameters are tunned on the development data split of the Wiki-TableQuestions table. We evaluate the model every 500 steps on the validation set, and choose the best performing model after reaching 50,000 training steps using the early stopping procedure. Each model variant is trained eight times and the best one of each variant is eventually run against the test set.

[2] In some cases it happens that, when executed on a particular table, a logical form gives the correct answer by chance without being the real translation of the input question.

Baselines

System	Accuracy
Neural Programmer	34.2%
Neural Programmer Ensemble (15 models)	37.7%
PL2015	37.1%

Our Models

System	Accuracy
CNN-DOTPRODUCT	31.4%
CNN-BILINEAR	20.4%
CNN-FC	30.4%
CNN-FC-BILINEAR	33.3%
RNN-FC-BILINEAR-SEP	29.6%
CNN-FC-BILINEAR-SEP (best single)	34.8%
Ensemble (15 models)	**38.7%**

Table 1: Results on the WikiTableQuestions dataset of different systems. Notation used for our models: <sentence embedding type> - <similarity network>.

4.3 Results

Table 1 shows the performance of our models compared to Neural Programmer [Neelakantan *et al.*, 2016] and PL2015 [Pasupat and Liang, 2015] baselines. The best performing single model is a linear combination between BILINEAR and FC models, namely CNN-FC-BILINEAR-SEP, that gives an accuracy of 34.8%. One explanation for this is that the two methods are able to recover different types of errors. Our best final model is an ensemble of 15 single models, reaching a state-of-the-art accuracy for this task of 38.7%. The score of the ensemble is calculated by averaging over the normalized scores of its constituents. The significant increase in performance of the ensemble over the single model shows that the different models learn unique features.

In an additional experiment, we use a recurrent neural network (RNN) for the sentence embedding, observing that the model RNN-FC-BILINEAR-SEP performs significantly worse than the corresponding CNN variant. RNNs are known to work well with sequence data, while CNNs can capture

Question	Paraphrase
Which association entered last?	association of last row
	association of row with highest number of joining year
What is the total of all the medals?	count all rows
	number of total of nation is total
How many episodes were originally aired before December 1965?	count original air date as date <= 12 1965
	count original air date as date <12 1965

Table 2: Example questions highlighting common mistakes our model makes. Correctness of a logical form is indicated by green color, whereas red color represents an incorrect logical form.

System	Accuracy (Dev)
CNN-FC-BILINEAR-SEP	34.1%
without Dropout	33.3%
without Character Embeddings	33.8%
without GloVe	32.4%
without Paraphrasing	33.1%

Table 3: Contributions of each component of our model.

System	Amount (%)
Lookup	10.8%
Aggregation + next, previous	39.8%
Superlatives	30.1%
Arithmetic and Comparisons	19.3%

Table 4: Type distribution of correctly answered questions.

patterns in a bag-of-n-grams manner, which is more suitable for the paraphrases produced by Algorithm 1.

There are a few reasons for the low accuracy obtained on this task by various methods (including ours) compared to other NLP problems. Weak supervision, small training size and a high percentage of unanswerable questions[3] contribute to this difficulty.

4.4 Ablation Studies

For a better understanding of our model, we investigate the usefulness of various components with an ablation study shown in Table 3. Leaving out the character embeddings has a marginal effect on accuracy. Regularizing the fully connected layers using dropout is important. However, the biggest impact on accuracy comes from using GloVe pre-trained word vectors to initialize the token embeddings, since switching to random initializations significantly decreases the accuracy.

In order to test the effect of the paraphrasing on the quality of the results, we conduct an additional experiment by replacing the paraphrase with the raw strings of the Lambda DCS expressions. The results are worse by a small margin, confirming that the paraphrasing method is not inducing additional errors. Moreover, our neural network component has the biggest impact in the success of our method.

4.5 Analysis of Correct Answers

To gain further insight into our approach, we analyze how well our best single model, CNN-FC-BILINEAR-SEP, performs on various question types. We manually annotate 80 randomly chosen questions that are correctly answered by our model. Results are shown in Table 4. The biggest contribution to accuracy stems from questions containing aggregation, next or previous operations[4], even though they only account to 20.5% of all the questions [5].

4.6 Error Analysis

The questions our models do not answer correctly can be split into two categories: either a correct logical form is not generated, or our scoring models do not rank the correct one at the top. In many cases, the correctness of a logical forms depends highly on the table structure. We perform a qualitative analysis presented in Table 2 to reveal common question types our models often rank incorrectly. The first two examples show questions whose correct logical form depends on the structure of the table. In these cases a bias towards the more general logical form is often exhibited. The third example shows that our model has difficulty distinguishing operands with slight modification (e.g. smaller and smaller equals), which may be due to weak-supervision. As we do not use or have access to the ground truth logical form during training, but only to the correct answer, queries using operands with slight modifications, would yield the same answer except in the edge case.

5 Conclusion

In this paper we propose a two stage QA system for semi-structured tables. The first stage consists of a standard method for generating candidate logical forms and a simple approach for transforming logical forms into textual paraphrases understandable by non-expert users. The second stage is a fully neural model that ranks the candidate logical forms indirectly through their respective paraphrases, eliminating the need for manually designed features. Experiments show that an ensemble of our models reaches state-of-the-art accuracy on the WikiTableQuestions dataset, thus indicating its capability to answer complex, multi-compositional questions. In the future we plan to advance this work by extending it to be able to reason on queries across multiple tables and work on an

[3] [Pasupat and Liang, 2015] state that 21% of questions cannot be answered because of various issues like annotation errors or tables requiring advanced normalization.

[4] Next/previous fetch the table row below/above the current.

[5] The distribution of the various question types in the WikiTablesQuestions dataset can be found at: `http://cs.stanford.edu/~ppasupat/resource/ACL2015-poster.pdf`

end-to-end approach where a joint training of both stages can be achieved.

Our code is publicly available at `https://github.com/dalab/neural_qa`.

References

[Abadi *et al.*, 2016] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, et al. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*, 2016.

[Berant and Liang, 2014] Jonathan Berant and Percy Liang. Semantic parsing via paraphrasing. In *ACL (1)*, pages 1415–1425, 2014.

[Bordes *et al.*, 2014a] Antoine Bordes, Sumit Chopra, and Jason Weston. Question answering with subgraph embeddings. *arXiv preprint arXiv:1406.3676*, 2014.

[Bordes *et al.*, 2014b] Antoine Bordes, Jason Weston, and Nicolas Usunier. Open question answering with weakly supervised embedding models. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 165–180. Springer, 2014.

[Clark *et al.*, 2016] Peter Clark, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Oyvind Tafjord, Peter D Turney, and Daniel Khashabi. Combining retrieval, statistics, and inference to answer elementary science questions. In *AAAI*, pages 2580–2586, 2016.

[Clevert *et al.*, 2015] Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. Fast and accurate deep network learning by exponential linear units (elus). *arXiv preprint arXiv:1511.07289*, 2015.

[Denton *et al.*, 2015] Emily Denton, Jason Weston, Manohar Paluri, Lubomir Bourdev, and Rob Fergus. User conditional hashtag prediction for images. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1731–1740. ACM, 2015.

[Dong *et al.*, 2015] Li Dong, Furu Wei, Ming Zhou, and Ke Xu. Question answering over freebase with multi-column convolutional neural networks. In *ACL (1)*, pages 260–269, 2015.

[Jauhar *et al.*, 2016] Sujay Kumar Jauhar, Peter D Turney, and Eduard Hovy. Tables as semi-structured knowledge for question answering. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, volume 1, pages 474–483, 2016.

[Kim *et al.*, 2015] Yoon Kim, Yacine Jernite, David Sontag, and Alexander M Rush. Character-aware neural language models. *arXiv preprint arXiv:1508.06615*, 2015.

[Kim, 2014] Yoon Kim. Convolutional neural networks for sentence classification. *arXiv preprint arXiv:1408.5882*, 2014.

[Kingma and Ba, 2014] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[Kwiatkowski *et al.*, 2013] Tom Kwiatkowski, Eunsol Choi, Yoav Artzi, and Luke Zettlemoyer. Scaling semantic parsers with on-the-fly ontology matching. In *In Proceedings of EMNLP. Percy*. Citeseer, 2013.

[Liang, 2013] Percy Liang. Lambda dependency-based compositional semantics. *arXiv preprint arXiv:1309.4408*, 2013.

[Mikolov *et al.*, 2013] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119, 2013.

[Neelakantan *et al.*, 2016] Arvind Neelakantan, Quoc V Le, Martin Abadi, Andrew McCallum, and Dario Amodei. Learning a natural language interface with neural programmer. *arXiv preprint arXiv:1611.08945*, 2016.

[Pasupat and Liang, 2015] Panupong Pasupat and Percy Liang. Compositional semantic parsing on semi-structured tables. *arXiv preprint arXiv:1508.00305*, 2015.

[Pennington *et al.*, 2014] Jeffrey Pennington, Richard Socher, and Christopher D Manning. Glove: Global vectors for word representation. In *EMNLP*, volume 14, pages 1532–1543, 2014.

[Reddy *et al.*, 2014] Siva Reddy, Mirella Lapata, and Mark Steedman. Large-scale semantic parsing without question-answer pairs. *Transactions of the Association for Computational Linguistics*, 2:377–392, 2014.

[Socher *et al.*, 2013] Richard Socher, Alex Perelygin, Jean Y Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, Christopher Potts, et al. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the conference on empirical methods in natural language processing (EMNLP)*, volume 1631, page 1642. Citeseer, 2013.

[Srivastava *et al.*, 2014] Nitish Srivastava, Geoffrey E Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(1):1929–1958, 2014.

[Wang *et al.*, 2015] Yushi Wang, Jonathan Berant, Percy Liang, et al. Building a semantic parser overnight. In *ACL (1)*, pages 1332–1342, 2015.

[Yin *et al.*, 2015] Pengcheng Yin, Zhengdong Lu, Hang Li, and Ben Kao. Neural enquirer: Learning to query tables with natural language. *arXiv preprint arXiv:1512.00965*, 2015.

[Zhang *et al.*, 2015] Xiang Zhang, Junbo Zhao, and Yann LeCun. Character-level convolutional networks for text classification. In *Advances in neural information processing systems*, pages 649–657, 2015.